

eBay’s Market Intermediation Problem

Kamal Jain*

Christopher A. Wilkens†

Abstract

We study the optimal mechanism design problem faced by a market intermediary who makes revenue by connecting buyers and sellers. We first show that the optimal intermediation protocol has substantial structure: it is the solution to an algorithmic pricing problem in which seller’s costs are replaced with virtual costs, and the sellers’ payments need only depend on the buyer’s behavior and not the buyer’s actual valuation function.

Since the underlying algorithmic pricing problem may be difficult to solve optimally, we study specific models of buyer behavior and give mechanisms with provable approximation guarantees. We show that offering only the single most profitable item for sale guarantees an $\Omega(\frac{1}{\log n})$ fraction of the optimal revenue when item value distributions are independent and have monotone hazard rates. We also give constant factor approximations when the buyer considers all items at once, k items at once, or items in sequence.

1 Introduction

The internet has connected people in ways unfathomable a few decades ago. This newfound ease of communication has ushered in a new era in commerce, yet, alone the internet is just a tool — without intermediaries like Google, eBay, and Amazon to help buyers navigate and connect with sellers, e-commerce would not have flourished. In particular, sites like eBay and Amazon have established themselves as nexuses connecting buyers and sellers in ways that were impossible without the internet. This raises a natural design question: given a group of buyers and sellers, what is the optimal way to intermediate between them? We call this the market intermediation problem.

Market intermediation is fundamental to e-commerce. First, it solves an informational problem by connecting buyers and sellers who were not already aware of each other. For example, without a central portal like Google or eBay, a buyer might only shop at the online stores for sellers he already knew about offline. Second, an intermediary reduces transaction costs by making the most appropriate connections between buyers and sellers. Even if buyers knew all the sellers on the internet, there is a small bargaining cost for a buyer and seller to exchange information about supply, demand, and prices. The wide variety of buyers and sellers on the internet would make such costs prohibitively large — an intermediary reduces costs by aggregating information, allowing buyers and sellers to be matched efficiently.

Our work has roots in the classic economic analysis of bilateral trade and double auctions. In 1983, Myerson and Satterthwaite [12] asked a simple question: will two self-interested people always

*eBay Research Labs, San Jose, CA.

†University of California at Berkeley, Berkeley, CA. This work was done while the author was an intern at eBay Research Labs.

trade when it is beneficial? The intriguing answer — there is no equilibrium in which two parties trade whenever it is beneficial without the help of an intermediary — derives from constraints on truthful payments, and their subsequent analysis of “Trading with a Broker” precisely models our scenario with a single buyer and a single seller. Myerson and Satterthwaite’s result [12] was the first in a line of research on double auctions (e.g. [10]), where many buyers and sellers place bids on a single commodity. More recent contributions from computer scientists have considered efficient algorithms and approximations (e.g. [15, 9, 8]) as well as the effect of a network structure [4]. The most relevant work to our own is a very recent paper by Deng et al. [8], who give approximation algorithms for optimizing revenue in Bayesian settings with multiple buyers, sellers, and items.

The market intermediation problem also has roots in the more recent computer science study of algorithmic pricing: given a collection of items and knowledge of buyer behavior, how should the items be priced? This basic computational question was first studied by Rusmevichientong et al. [13] and Aggarwal et al. [1] in a rank-buying model. Subsequent research considered a wide variety of other models, the most relevant results being Chawla, et al. [6, 7] who gave approximations for the case of a single unit-demand buyer and Cai and Daskalakis [5] who gave an FPTAS for the same setting.

The market intermediation problem combines these two research areas: interaction with the buyer is essentially an algorithmic pricing problem, yet the items for sale are brought by selfish agents as in double auctions.

Results. We study the market intermediation problem in a setting with a single unit-demand buyer and a group of n sellers. First, we show that the optimal solution to this problem has substantial structure — it is monotone for the sellers (necessary and sufficient for incentive compatibility), buyer prices can be reduced to an algorithmic pricing problem in which the intermediary is itself the seller, and the payments to the sellers need only depend on the costs and likelihood of sale (not the precise value of the buyer).

The market intermediation problem thus reduces to a computational one: the intermediary must solve an algorithmic pricing problem to choose the prices the buyer sees and then compute payments to the sellers. Since the sellers are single-parameter agents, their payments are given by a straightforward integral expression [11, 2], and even if that integral is impossible to compute we can use tricks like Archer et al. [3] to maintain truthfulness in expectation. On the other hand, the optimal prices for the buyer may be more difficult to compute, so we study a variety of efficient approximation strategies for the intermediary.

The first approximation we consider is the simplest strategy for the intermediary: pick the single best item and sell it. We show that this guarantees an $\Omega(\frac{1}{\log n})$ fraction of the optimal revenue when the items’ values are drawn from independent, monotone hazard rate distributions.

Next, we consider approximations when the buyer picks his favorite item from a set. When the buyer considers all items, we adapt the anonymous virtual reserve price of Chawla et al. [7] to give a $\frac{1}{2}$ -approximation to the optimal revenue.¹ In the more realistic setting where the intermediary can show the buyer at most k items, the intermediary must choose which k items to show and then compute prices. Here, the intermediary wants to approximate the optimal revenue that can be obtained by selling at most k items. We show that an anonymous virtual reserve price can

¹A result of Deng et al. [8] implies a weaker $\frac{1}{4}$ -approximation in this setting. Their result is technically incomparable to ours — they study a more general setting when there are multiple buyers who each demand an arbitrary number of goods, but they make substantially stronger assumptions about the representation of probability distributions.

be combined with a greedy algorithm to give a $\frac{e-1}{2e}$ -approximation, provided the intermediary can estimate the expected maximum value of a set of items (results of Cai and Daskalakis [5] show that this requirement is nontrivial, as we discuss later).

Finally, we consider a buyer who looks at items in order until he finds one he is willing to buy. We show that the order is less important than the prices: choosing any ordering and computing the optimal prices for that order gives a $\frac{1}{2}$ -approximation to the revenue of the optimal cascade.

2 The Market Intermediation Problem

The market intermediation problem models an intermediary connecting a buyer to a group of n sellers. Each seller i is offering a distinct, indivisible item and faces a privately-known cost of $c_i \sim \mathcal{C}_i$ when a sale is made (e.g. the on-demand manufacturing cost). Sellers report their costs when they come to the mechanism. We use c to denote the vector of costs and $\mathcal{C} = \mathcal{C}_1 \times \dots \times \mathcal{C}_n$ to denote the distribution of c . We assume that the distributions for costs are independent and have monotone hazard rates² (see Section 2.1).

A unit-demand buyer with type $t \in T$ (where T is independent of \mathcal{C}) comes to the intermediary to buy an item. After interacting with the intermediary, the buyer purchases item i with probability $x_i(t, c)$, paying an (expected) price $p_i^B(t, c)$.

The intermediary otherwise has complete control over the interaction between the buyer and the seller — its goal is to maximize revenue through a truthful protocol, that is, to pay sellers in such a way that revealing their true costs c is a dominant strategy.

The Independent Values Setting. Our algorithmic results assume the buyer’s valuation is described by n values v_i drawn from independent distributions \mathcal{V}_i .

The buyer is generally presented with a set of prices p_i^B and chooses the item that maximizes $v_i - p_i^B$; however, he may not consider all items available. For example, a buyer on eBay may only consider the first page of search results or may look at items until he finds something he likes. We discuss specific attention models in Sections 5 and 6.

2.1 Distributions, Hazard Rates, and Virtual Values

We represent distributions with density functions $f(z)$ and cumulative density functions $F(z)$. For such a distribution $F(z)$, the *hazard rate* is given by

$$h(z) = \frac{f(z)}{1 - F(z)} .$$

The following identity directly relates $F(z)$ to the hazard rate:

$$F(z) = 1 - e^{-\int_0^z h(x)dx} . \tag{1}$$

A noteworthy distribution is the exponential distribution $E_\gamma(z) = 1 - e^{-\gamma z}$, which has a constant hazard rate $h_{E_\gamma}(z) = \gamma$.

²The monotone hazard rate requirement is applied to the distribution $-\mathcal{C}_i$ so as to guarantee that the virtual cost $\theta_{\mathcal{C}_i}(c_i)$ is monotone.

A distribution F is said to have a *monotone hazard rate (MHR)* if $h(z)$ is non-decreasing. MHR distributions are convenient because the Myersonian *virtual value* [11] is monotone. Myerson's virtual value $\phi_F(v)$ represents the marginal revenue obtained by selling to a buyer with value $v \sim F(v)$:

$$\phi_F(v) = v - \frac{1}{h(v)} = v - \frac{1 - F(v)}{f(v)} .$$

For this work, we also define the *virtual cost* of a seller as

$$\theta_F(c) = c + \frac{F(c)}{f(c)} .$$

This represents the marginal cost to the intermediary when a seller with cost c makes a sale.

2.2 Optimal Single-Item Mechanisms

We highlight two important classical optimality results. First, Myerson's optimal auction [11] says that the optimal way to sell a single item to a single buyer with an MHR distribution is to choose the price p_M satisfying $\phi_V(p_M) = 0$. When the seller has an outside option \mathcal{R}_{OUT} and faces a cost c , Myerson's result says that the optimal price η satisfies

$$\eta - c - \mathcal{R}_{OUT} = \frac{1}{h_V(\eta)} .$$

Myerson and Satterthwaite [12] later considered the case where a selfish player holds the item (instead of the auctioneer). Their result precisely corresponds to the market intermediation problem with one seller ($n = 1$):

Theorem 2.1 (Myerson and Satterthwaite) *In the revenue optimal intermediation with a single seller and MHR distributions, the buyer purchases the item if and only if $\phi_V(v) - \theta_C(c) \geq 0$.*

3 Revenue Optimization in the General Setting

In this section, we show that the revenue optimal intermediation scheme has three nontrivial features:

1. *The revenue optimal scheme is monotone for the sellers.* Following classical results for single-parameter agents [11, 2], this is a necessary and sufficient criterion for seller incentive compatibility.
2. *The prices seen by the bidder p_i^B can be computed as if the intermediary were the seller and faced a cost of $\theta_{C_i}(c_i)$.* This gives a natural reduction from the market intermediation problem to an algorithmic pricing problem.
3. *The payments to the sellers p_i^S only depend on the costs and purchase decision of the buyer, not on the exact values v .* This implies that the intermediary need not ask the buyer for his actual value.

Algorithm 1 gives a general mechanism for the market intermediation problem incorporating these features.

We first prove seller monotonicity:

ALGORITHM 1: A general mechanism for the market intermediation problem.

- 1 **Solicit** costs c_i from sellers;
 - 2 **Compute** virtual costs $\theta_{c_i}(c_i)$;
 - 3 **Compute** a protocol for interacting with the buyer, including what items are shown and at what prices; // The protocol may be restricted when the buyer's attention is limited.
 - 4 **Pay** the seller whose item was bought $p_i^S(c)$; // p_i^S need not depend on the precise value of the buyer, only on the likelihood of sale.
-

Theorem 3.1 *The revenue optimal instantiation of the general algorithm (Algorithm 1) is monotone for the sellers, i.e. the probability of selling item i is decreasing in c_i .*

Proof: Suppose T is represented as a probability measure. Define the cost vector c^{+i} as the vector that increases i 's cost by δ , i.e.

$$c_j^{+i} = \begin{cases} c_j + \delta & i = j \\ c_j & \text{otherwise.} \end{cases}$$

Consider the optimal intermediation scheme for a type t and cost vector c . Let $x_i(t, c)$ denote the probability that the buyer purchases item i and $p_i^B(t, c)$ denote the expected price for this scheme as defined in Section 2.

Since the distribution T is independent of the cost distributions, the buyer's behavior does not depend on c . Thus, when the types and costs are t and c respectively, optimality implies that the intermediary makes more money selling according to $x_i(t, c)$ and $p_i^B(t, c)$ than if it acted as if the cost vector was c^{+j} . Thus:

$$\sum_i \int x_i(t, c)(p_i^B(t, c) - \theta_{c_i}(c_i))dT \geq \sum_i \int x_i(t, c^{+j})(p_i^B(t, c^{+j}) - \theta_{c_i}(c_i))dT .$$

Likewise, optimality for c^{+j} implies

$$\sum_i \int x_i(t, c^{+j})(p_i^B(t, c^{+j}) - \theta_{c_i}(c_i^{+j}))dT \geq \sum_i \int x_i(t, c)(p_i^B(t, c) - \theta_{c_i}(c_i^{+j}))dT .$$

Combining these two and canceling terms gives

$$\int x_j(t, c)(\theta_{c_j}(c_j + \delta) - \theta_{c_j}(c_j))dT \geq \int x_j(t, c)(\theta_{c_j}(c_j + \delta) - \theta_{c_j}(c_j))dT .$$

Since the virtual costs are assumed to be monotone, this implies

$$\int x_j(t, c)dT \geq \int x_i(t, c)dT$$

as desired. ■

Given monotonicity, we can apply classical results to compute the revenue and show that the intermediary's pricing problem can be formulated as an algorithmic pricing problem with virtual costs:

Theorem 3.2 *The revenue-optimal intermediation protocol can be described as follows: for each cost vector c , solve the underlying algorithmic pricing problem for p assuming the cost of item i is $\theta_{C_i}(c_i)$.*

Proof: Let $x_i(t, c)$ be the probability that the buyer purchases item i and let π_t be the probability that it has type t . From [11, 2] we know that for incentive compatibility the payment of seller i must be

$$p_i^S(t, c) = -c_i x(t, c) - \int_0^{c_i} x(t, c_{-i}u) du .$$

We can thus write the expected revenue of the mechanism as

$$\mathbf{E}[\mathcal{R}] = \int f_C(c) \int \sum_i (p_i^B(c) x_i(t, c) + p_i^S(t, c)) dT dc .$$

Rearranging in the standard way gives

$$\mathbf{E}[\mathcal{R}] = \int f_C(c) \int \sum_i x_i(t, c) (p_i^B(c) - \theta_{C_i}(c_i)) dT dc .$$

That is, the revenue of the intermediary is precisely as if he faces a cost of $\theta_{C_i}(c_i)$ for item i . ■

Corollary 3.3 *An α approximation to the underlying algorithmic pricing problem without costs gives an α approximation to the market intermediation problem as long as it is monotone from the sellers' perspectives.*

Proof: Given the true cost vector c (incentive compatibility is guaranteed by monotonicity), the virtual costs $\theta_{C_i}(c_i)$ are constants. For each item i , we solve the algorithmic pricing problem with the distribution \mathcal{V}_i shifted by $\theta_{C_i}(c_i)$. This guarantees an α approximation for every c and thus an α approximation to the market intermediation problem. ■

Finally, we show that the intermediary need not explicitly query the buyer's value:

Theorem 3.4 *The revenue optimal instantiation of the general algorithm can be implemented so that seller payments depend only on costs and which item is bought, not the precise value of the buyer.*

Proof: Suppose truthful payments $p_i^S(t, c)$ exist. Charge $p_i^S(c) = \mathbf{E}_{t \sim T}[p_i^S(t, c)]$. ■

4 Selling the Best Single Item

A simple strategy for the intermediary is to find the single most profitable item and only show that item to the buyer. This strategy can generally guarantee at most a $\frac{1}{n}$ fraction of the optimal revenue; however, we show that this gives a much better approximation when the value of each item is independent and drawn from a MHR distribution \mathcal{V}_i :

Theorem 4.1 *When the buyer's valuations \mathcal{V}_i are independent and MHR, selling the single most profitable item (based on virtual costs $\theta_{C_i}(c_i)$) in an optimal auction generates at least a $\frac{1}{e \ln n}$ fraction of the optimal revenue in either the full attention or cascade model.*

To prove Theorem 4.1, we first to show that the benefit from selling an extra item with distribution \mathcal{V} in the cascade model is maximized when \mathcal{V} is an exponential distribution. We then bound the approximation for items with exponential distributions.

Suppose a seller wishes to sell a single item to a buyer whose value is drawn from the distribution \mathcal{V} (the seller's cost is 0). Let \mathcal{R}_M be the revenue generated by selling the item by itself using Myerson's optimal auction, and let p_M be the price at which Myerson's auction sells the item.

Suppose also that the seller has an outside option guaranteeing revenue of \mathcal{R}_{OUT} if the item does not sell. Let $\mathcal{R}_\mathcal{V} > \mathcal{R}_M$ be the revenue generated by the optimal auction when the seller considers his outside option, and let $\Delta\mathcal{R}_\mathcal{V} = \mathcal{R}_\mathcal{V} - \mathcal{R}_{OUT}$ be the optimal incremental revenue, i.e. the extra benefit to the seller on top of his outside option. Then we have the following lemma:

Lemma 4.2 *For a fixed value of \mathcal{R}_M , the incremental revenue $\Delta\mathcal{R}_\mathcal{V}$ is maximized when \mathcal{V} is an exponential distribution for any value of \mathcal{R}_{OUT} . In particular,*

$$\Delta\mathcal{R}_\mathcal{V} \leq \frac{1}{e\gamma} e^{-\gamma\mathcal{R}_{OUT}}$$

where $\gamma = \frac{1}{e\mathcal{R}_M}$.

Proof: First, we identify three distributions:

- \mathcal{V} is the true distribution of v as defined in the problem.
- E is the exponential distribution that generates the same revenue as \mathcal{V} in Myerson's optimal auction. We will see that this is the exponential distribution with $\gamma = \frac{1}{e\mathcal{R}_M}$.
- $\tilde{\mathcal{V}}$ is a hybrid defined through its hazard rate:

$$h_{\tilde{\mathcal{V}}} = \begin{cases} h_{\mathcal{V}}(v), & v \leq p_M \\ h_{\mathcal{V}}(p_M), & \text{otherwise.} \end{cases}$$

(The fact that $\tilde{\mathcal{V}}$ is well-defined follows from identity (1).)

Our proof shows that $\Delta\mathcal{R}_\mathcal{V} \leq \Delta\mathcal{R}_{\tilde{\mathcal{V}}} \leq \Delta\mathcal{R}_E$.

Preliminaries. For a distribution F , let η_F denote the optimal selling price if we account for the seller's outside option. Thus, η_F must satisfy

$$\eta_F - \mathcal{R}_{OUT} = \frac{1}{h_F(\eta_F)} .$$

The total (optimal) revenue \mathcal{R}_F generated by F is therefore

$$\mathcal{R}_F = (1 - F(\eta_F))\eta_F + F(\eta_F)\mathcal{R}_{OUT} .$$

Thus, the incremental revenue of selling at a fixed price p is

$$\begin{aligned} \mathcal{R}_F(p) - \mathcal{R}_{OUT} &= (1 - F(p))p + F(p)\mathcal{R}_{OUT} - \mathcal{R}_{OUT} \\ &= (1 - F(p))(p - \mathcal{R}_{OUT}) , \end{aligned}$$

and the optimal incremental revenue $\Delta\mathcal{R}_F$ is

$$\Delta\mathcal{R}_F = (1 - F(\eta_F))(\eta_F - \mathcal{R}_{OUT}) = \frac{1 - F(\eta_F)}{h_F(\eta_F)} = \mathcal{R}_M \frac{\frac{1-F(\eta_F)}{1-F(p_M)}}{\frac{h_F(\eta_F)}{h_F(p_M)}} .$$

Step 1: $\Delta\mathcal{R}_\mathcal{V} \leq \Delta\mathcal{R}_{\tilde{\mathcal{V}}}$.

Let z satisfy

$$1 - F_{\tilde{\mathcal{V}}}(z) = 1 - F_{\mathcal{V}}(\eta_{\mathcal{V}}) .$$

Since $h_{\mathcal{V}}(v)$ is assumed to be nondecreasing, the identity (1) implies

$$1 - F_{\tilde{\mathcal{V}}}(v) = e^{-\int_0^v h_{\tilde{\mathcal{V}}}(u)du} \geq e^{-\int_0^v h_{\mathcal{V}}(u)du} = 1 - F_{\mathcal{V}}(v) ,$$

and thus $z \geq \eta_{\mathcal{V}}$ because F is nondecreasing. We know that the optimal revenue $\mathcal{R}_{\tilde{\mathcal{V}}}$ is at least the revenue generated by selling at price z , so

$$\begin{aligned} \Delta\mathcal{R}_{\tilde{\mathcal{V}}} &\geq (1 - F_{\tilde{\mathcal{V}}}(z))(z - \mathcal{R}_{OUT}) \\ &\geq (1 - F_{\mathcal{V}}(\eta_{\mathcal{V}}))(\eta_{\mathcal{V}} - \mathcal{R}_{OUT}) = \Delta\mathcal{R}_{\mathcal{V}} . \end{aligned}$$

Thus, the incremental revenue from $\tilde{\mathcal{V}}$ is at least as large as the incremental revenue from \mathcal{V} .

Step 2: $\Delta\mathcal{R}_{\tilde{\mathcal{V}}} \leq \Delta\mathcal{R}_E$.

As shown above, the incremental revenue $\Delta\mathcal{R}_{\tilde{\mathcal{V}}}$ is

$$\Delta\mathcal{R}_{\tilde{\mathcal{V}}} = \mathcal{R}_M \frac{\frac{1-F_{\tilde{\mathcal{V}}}(\eta_{\tilde{\mathcal{V}}})}{1-F_{\tilde{\mathcal{V}}}(p_M)}}{\frac{h_{\tilde{\mathcal{V}}}(\eta_{\tilde{\mathcal{V}}})}{h_{\tilde{\mathcal{V}}}(p_M)}} = \mathcal{R}_M \frac{e^{-(\eta_{\tilde{\mathcal{V}}}-p_M)h_{\mathcal{V}}(p_M)}}{1} .$$

By the optimality of $\eta_{\tilde{\mathcal{V}}}$ and p_M as well as the definition of $\tilde{\mathcal{V}}$ we know that

$$\eta_{\tilde{\mathcal{V}}} - \mathcal{R}_{OUT} = \frac{1}{h_{\tilde{\mathcal{V}}}(\eta_{\tilde{\mathcal{V}}})} = \frac{1}{h_{\mathcal{V}}(p_M)} = p_M ,$$

and so $\eta_{\tilde{\mathcal{V}}} - p_M = \mathcal{R}_{OUT}$. Thus,

$$\Delta\mathcal{R}_{\tilde{\mathcal{V}}} = \mathcal{R}_M e^{-\mathcal{R}_{OUT}h_{\mathcal{V}}(p_M)} .$$

Similar analysis of the exponential distribution E shows that

$$\Delta\mathcal{R}_E = \mathcal{R}_M e^{-\mathcal{R}_{OUT}\gamma} .$$

Thus, to show $\Delta\mathcal{R}_E \geq \Delta\mathcal{R}_{\tilde{\mathcal{V}}}$ we need only show $\gamma \leq h_{\mathcal{V}}(p_M)$.

One can check that the optimal revenue of E in the absence of \mathcal{R}_{OUT} is

$$\frac{1 - F_E(\frac{1}{\gamma})}{h_E(\frac{1}{\gamma})} = \frac{1}{e\gamma} .$$

Thus, by definition of E we have chosen γ so that $\mathcal{R}_M = \frac{1}{e\gamma}$. Using the facts that $\mathcal{R}_M = \frac{1-F_V(p_M)}{h_V(p_M)}$, h_V is nondecreasing, and $p_M = \frac{1}{h_V(p_M)}$ we get

$$\begin{aligned} h_V(p_M) &= e\gamma(1 - F_V(p_M)) \\ &= e\gamma e^{-\int_0^{p_M} h_V(v) dv} \\ &\geq e\gamma e^{-p_M h_V(p_M)} = e\gamma e^{-1} = \gamma \end{aligned}$$

This gives

$$\Delta\mathcal{R}_{\hat{V}} = \mathcal{R}_M e^{-\mathcal{R}_{OUT} h_V(p_M)} \leq \mathcal{R}_M e^{-\mathcal{R}_{OUT} \gamma} = \Delta\mathcal{R}_E$$

as desired. ■

We can now prove our approximation theorem:

Proof: (of Theorem 4.1) Following Corollary 3.3, it is sufficient to consider the pricing problem of a single seller without costs. Let Π_{OPT} be the optimal ordering of items in a cascade. Let \mathcal{R}_k be the revenue if we only sold the last k items, and let $\Delta\mathcal{R}_k = \mathcal{R}_k - \mathcal{R}_{k-1}$ be the incremental benefit of being able to sell the k th item. Then by Lemma 4.2 we have

$$\mathcal{R}_k = \mathcal{R}_{k-1} + \Delta\mathcal{R}_k \leq \mathcal{R}_{k-1} + \frac{1}{e\gamma_k} e^{-\gamma_k \mathcal{R}_{k-1}} \leq \mathcal{R}_{k-1} + \frac{1}{e\gamma_{min}} e^{-\gamma_{min} \mathcal{R}_{k-1}}$$

where $\frac{1}{e\gamma_k}$ is the Myersonian revenue of selling item i and γ_{min} is the minimum over all k . We use induction to prove that $\mathcal{R}_k \leq \frac{\ln k}{\gamma_{min}}$.

Base case $k = 2$:

$$\mathcal{R}_k \leq \frac{1}{e\gamma_{min}} e^{-\frac{1}{e}} < \frac{\ln 2}{e\gamma_{min}}$$

Inductive step:

$$\begin{aligned} \mathcal{R}_k &\leq \mathcal{R}_{k-1} + \frac{1}{e\gamma_{min}} e^{-\gamma_{min} \mathcal{R}_{k-1}} \\ &= \frac{\ln(k-1)}{\gamma_{min}} + \frac{1}{e\gamma_{min}} e^{-\gamma_{min} \frac{\ln(k-1)}{\gamma_{min}}} \\ &\leq \frac{1}{\gamma_{min}} \left(\ln(k-1) + \ln\left(\frac{k}{k-1}\right) \right) \\ &= \frac{\ln k}{\gamma_{min}} \end{aligned}$$

Thus, we have

$$\mathcal{R}_n \leq (e \ln n) \max_i \mathcal{R}_M^i$$

as desired.

Since the optimal revenue from a cascade is an upper bound on the optimal revenue in a full attention model, this implies a $\frac{1}{e \ln n}$ approximation for the full attention setting as well. ■

5 Approximation with (Limited) Full Attention

As noted earlier, buyers often have limited attention — a buyer might only look at the first page of search results or may scan down a list until he finds an item he is willing to buy. In this section we use anonymous virtual reserve prices to give approximations when the buyer considers a set of items and picks his favorite. We first consider *full attention model*, where the buyer is presented with a set of prices p_i^B and buyer considers all possible items, then we consider the more realistic *k-limited attention model* where the buyer only looks at k items. In the later case, the intermediary must choose which k items to show the buyer.

5.1 Full Attention

Chawla et al. [7] show how to compute an anonymous virtual reserve price that guarantees a $\frac{1}{2}$ -approximation to the algorithmic pricing problem — we show that this gives a $\frac{1}{2}$ -approximation in the full attention model.

The full attention model is effectively a special case of a double auction. In their very recent work, Deng et al. [8] give a $\frac{1}{4}$ -approximation in the more general setting with multiple buyers, each of whom are interested in buying an arbitrary (finite) number of items. This gives a (weaker) $\frac{1}{4}$ -approximation in our setting; however, their work also makes substantially stronger assumptions about the auctioneer’s access to the buyer’s distributions. As such, our result is substantially stronger for our single-buyer setting.

Theorem 5.1 *The $\frac{1}{2}$ -approximation of Chawla et al. [7] for the algorithmic pricing problem with a single unit demand bidder gives a $\frac{1}{2}$ -approximation to the market intermediation problem.*

Proof: Following Corollary 3.3, we need only show that the approximation of Chawla et al. [7] is monotone, i.e. the likelihood of the seller purchasing item i is decreasing in c_i . Their algorithm applied to our setting can be described as follows:

1. Let $z = \max_i \phi_{\mathcal{V}_i}(v_i)$, and let \mathcal{Z} be the distribution of z when $v_i \sim \mathcal{V}_i$.
2. Let r be the median³ of the distribution \mathcal{Z} .
3. Set $p_i^B = \phi^{-1}(r)$

Let $s_i = v_i - p_i$ be the surplus if the bidder buys item i . Note that the bidder’s decision depends only on the surpluses and not on the specific values v_i and p_i .

Let \mathcal{V}_i^δ denote the distribution \mathcal{V}_i shifted right by δ . Let p be the price vector given by computing this approximation with \mathcal{V}_i , and let p^δ be the price vector computed from \mathcal{V}_i^δ for all i .

We first show that $p_i^\delta = p_i + \delta$ and that, consequently, the probability the buyer buys any given item is the same. By definition of ϕ , the distribution of $\phi_{\mathcal{V}_i}(v_i)$ shifts by δ :

$$\phi_{\mathcal{V}_i^\delta}(v_i + \delta) = v_i + \delta - \frac{1}{h_{\mathcal{V}_i}(v_i)} = \phi_{\mathcal{V}_i}(v_i) + \delta .$$

³Chawla, Hartline, Malec, and Sivan [7] study a more general setting in which one cannot simply take r to be the median; however, the median suffices for our simpler setting [6, 14].

It follows that the median of the max distribution must also shift by δ , so $r^\delta = r + \delta$. Now p_i^δ is chosen to satisfy

$$\phi_{\mathcal{V}_i^\delta}(p_i^\delta) = r^\delta = r + \delta .$$

Notice that

$$\phi_{\mathcal{V}_i^\delta}(p_i + \delta) = \phi_{\mathcal{V}_i}(p_i) + \delta = r + \delta ,$$

so $p_i^\delta = p_i + \delta$. Thus, both v_i and p_i increased by δ , so the surpluses remain the same and the probability that the buyer chooses i is the same under \mathcal{V}_i as it is under \mathcal{V}_i^δ .

Next, consider shifting only the valuation of item i and selling at the price vector $(p_{-i}, p_i + \delta)$. Again, the surpluses remain the same, so the probability that the buyer chooses i is the same as it was before the shift. However, since the value r increases every time we shift a distribution to the right, for items $j \neq i$ we know that p_j is lower than the price given by the Chawla et al. approximation. Similarly, we know that $p_i + \delta$ is higher than the price at which the approximation would sell item i . Thus, in every outcome the surplus s_i is greater under the approximation than it is under $(p_{-i}, p_i + \delta)$ while the surpluses s_j are smaller. It follows that the likelihood the buyer chooses item i can only increase, as desired. ■

5.2 k -Limited Attention

In the k -limited attention model, the intermediary can only show k items to the user (e.g. the first page of search results). The buyer subsequently chooses the item maximizing $v_i - p_i^B$. We show that greedily choosing the set of items to show gives a constant-factor approximation.

Implementing the greedy algorithm requires estimating $\mathbf{E}[\max_{i \in S} v_i]$ for an arbitrary set S . Doing so efficiently is nontrivial — results of Cai and Daskalakis [5] show that the distribution of $\max_{i \in S} v_i$ may not be sufficiently concentrated to estimate $\mathbf{E}[\max_{i \in S} v_i]$ from samples for arbitrary distributions. We leave the problem of estimating $\mathbf{E}[\max_{i \in S} v_i]$ as an open question.

Conditioned on the intermediary's ability to estimate $\mathbf{E}[\max_{i \in S} v_i]$, the greedy algorithm can be combined with the anonymous virtual reserve construction of Chawla et al. [7] to get a constant-factor approximation:

Theorem 5.2 *Assuming the auctioneer can efficiently estimate $\mathbf{E}_{v_i \sim \mathcal{V}_i}[\max_{i \in S} \phi_{\mathcal{V}_i}(v_i)]$ for any set S , using a greedy algorithm to select which k items to show and then using the approximation of Chawla et al. [7] to compute prices gives an $\frac{e-1}{2e}$ approximation in the k -limited attention model.*

Proof: The approximation of Chawla et al. [7] used in Theorem 5.1 guarantees a $\frac{1}{2}$ -approximation to $\mathbf{E}_{v_i \sim \mathcal{V}_i}[\max_{i \in S} \phi_{\mathcal{V}_i}(v_i)]$, which is the Myersonian revenue if the buyer were split into n competing buyers (one for each good). This is itself an upper bound on the optimal revenue in the single-buyer setting.

Thus, if we knew the set S of items (where $|S| = k$) that maximized the expected virtual value $\mathbf{E}_{v_i \sim \mathcal{V}_i}[\max_{i \in S} \phi_{\mathcal{V}_i}(v_i)]$, then Theorem 5.1 would imply a $\frac{1}{2}$ -approximation here. Instead, we approximate S using the greedy algorithm. As shown in Lemma 5.4 (below), the greedy algorithm gives a set S whose expected virtual value is an $\frac{e-1}{e}$ -approximation to the optimal S , giving an $\frac{e-1}{2e}$ -approximation to the optimal revenue. The greedy algorithm is monotone from the sellers' perspective, proving the theorem. ■

It remains to prove that the greedy algorithm gives an $\frac{e-1}{e}$ -approximation as claimed in Theorem 5.2. This will follow from submodularity.

Let X_1, \dots, X_n be independent random variables, and let $S \subseteq [n]$ represent a subset of the X_i variables. Then:

Lemma 5.3 *The function $f(S) = \mathbf{E}[\max_{i \in S} X_i]$ is a monotone submodular set function.*

Proof: The function $f(S) = \max_{i \in S} X_i$ is a monotone submodular function, therefore by linearity of expectation so is $f(S) = \mathbf{E}[\max_{i \in S} X_i]$. ■

The following lemma follows from the standard analysis of the maximum coverage problem:

Lemma 5.4 *Suppose we asked to select a set S of size k that maximizes $\mathbf{E}[\max_{i \in S} X_i]$. The greedy algorithm gives a $(1 - \frac{1}{e})$ approximation for this problem.*

6 Approximation with Cascading Attention

In the cascading attention model, the intermediary has a set of slots. The buyer looks at slots according to an order Π and buys the first item for which $v_i \geq p_i$. Knowing this, the intermediary decides the matching of items to slots and the prices p_i^B for each item

Theorem 6.1 *Given any fixed ordering Π of items, the optimal prices for that ordering generate at least $\frac{1}{2}$ the revenue of the optimal cascade.*

Proof: Define the cumulative revenue \mathcal{R}_k under ordering Π as the revenue if we only sold the last k items, and the incremental revenue $\Delta\mathcal{R}_k$ is $\Delta\mathcal{R}_k = \mathcal{R}_k - \mathcal{R}_{k-1}$.

First, we observe that the optimal prices for a given ordering Π can be computed greedily. Because distributions are independent, the optimal price for item k is simply the optimal single-item auction price given an outside option of \mathcal{R}_{k-1} . Thus, the optimal prices can be computed recursively starting with the last item in Π . Lemma 6.2 (below) therefore implies that the incremental revenue of an item is decreasing in its cumulative revenue.

Now, let z be the smallest value of k such that $\mathcal{R}_k > \frac{1}{2}\mathcal{R}^{OPT}$ in the optimal order Π^{OPT} . If the cumulative revenue of this item in Π is at least $\frac{1}{2}\mathcal{R}^{OPT}$, then we are done.

Otherwise, the cumulative revenue of each item $k \geq z$ in Π^{OPT} is smaller in Π than it is in the optimal cascade Π^{OPT} . This implies that the incremental revenue $\Delta\mathcal{R}_k$ is greater in Π than it is in the optimum. However, the total incremental revenue of items $k \geq z$ is at least $\frac{1}{2}\mathcal{R}^{OPT}$ by construction, so the total revenue of Π must be at least $\frac{1}{2}\mathcal{R}^{OPT}$. ■

Lemma 6.2 *For an optimal single-item auction, the revenue from selling the item $\eta_{\mathcal{R}_{OUT}}(1 - F(\eta_{\mathcal{R}_{OUT}}))$ is decreasing in both the outside option \mathcal{R}_{OUT} and the total revenue of the auction $\mathcal{R}_{OUT} + \eta_{\mathcal{R}_{OUT}}(1 - F(\eta_{\mathcal{R}_{OUT}}))$.*

Proof: Optimality of $\eta_{\mathcal{R}_{OUT}}$ has a few relevant consequences. First, the optimal selling price $\eta_{\mathcal{R}_{OUT}}$ is increasing in \mathcal{R}_{OUT} . Second, combined with monotonicity of $\eta_{\mathcal{R}_{OUT}}$, optimality of $\eta_{\mathcal{R}_{OUT}}$ implies that the revenue from selling the item $\eta_{\mathcal{R}_{OUT}}(1 - F(\eta_{\mathcal{R}_{OUT}}))$ is also decreasing in \mathcal{R}_{OUT} . Finally, the total revenue is increasing in \mathcal{R}_{OUT} , so the revenue of selling the item $\eta_{\mathcal{R}_{OUT}}(1 - F(\eta_{\mathcal{R}_{OUT}}))$ is decreasing in the total revenue. ■

7 Conclusion and Open Questions

The market intermediation problem models a fundamental situation faced by companies like eBay, Amazon, and Google. Much work in economics and algorithmic game theory is relevant, but little directly addresses the problems these companies face. Our results show that the intermediary’s problem has significant structure and offers strategies with provable performance guarantees; however, it still leaves many open questions:

1. Which existing algorithmic pricing solutions are monotone from the sellers’ perspective?
2. What other value and attention models permit good approximations?
3. Are there ways the intermediary can leverage mechanisms that do not require direct revelation by the sellers, such as eBay’s commission model?
4. How does the model change with multiple bidders?
5. In practice, what is the best intermediation protocol for a company like eBay?

8 Thanks

We would like to thank Costis Daskalakis, Darrell Hoy, Christos Papadimitriou, George Pierrakos, Balasubramanian Sivan, and Steve Tadelis for many helpful discussions.

References

- [1] Gagan Aggarwal, Toms Feder, Rajeev Motwani, and An Zhu. Algorithms for multi-product pricing. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*, pages 72–83, 2004.
- [2] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *FOCS ’01: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*, page 482, Washington, DC, USA, 2001. IEEE Computer Society.
- [3] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *SODA ’03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 205–214, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.
- [4] Larry Blume, David Easley, Jon Kleinberg, and Eva Tardos. Trading networks with price-setting agents. In *Proceedings of the 8th ACM conference on Electronic commerce, EC ’07*, pages 143–151, New York, NY, USA, 2007. ACM.
- [5] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS ’11*, pages 522–531, Washington, DC, USA, 2011. IEEE Computer Society.

- [6] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM conference on Electronic commerce*, EC '07, pages 243–251, New York, NY, USA, 2007. ACM.
- [7] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 311–320, New York, NY, USA, 2010. ACM.
- [8] Xiaotie Deng, Paul Goldberg, Bo Tang, and Jinshan Zhang. Revenue maximization in a bayesian double auction market. In *ISAAC*, 2012.
- [9] Kaustubh Deshmukh, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Truthful and competitive double auctions. In *Proceedings of the 10th Annual European Symposium on Algorithms*, ESA '02, pages 361–373, London, UK, UK, 2002. Springer-Verlag.
- [10] D. Friedman and J. Rust. *The double auction market: institutions, theories, and evidence*. Proceedings volume in the Santa Fe Institute studies in the sciences of complexity. Addison-Wesley Pub. Co., 1993.
- [11] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [12] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, April 1983.
- [13] Paat Rusmevichientong, Benjamin Van Roy, and Peter W. Glynn. A nonparametric approach to multiproduct pricing. *Oper. Res.*, 54(1):82–98, January 2006.
- [14] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent non-negative random variables. *The Annals of Probability*, 12(4):pp. 1213–1216, 1984.
- [15] Peter R. Wurman, William E. Walsh, and Michael P. Wellman. Flexible double auctions for electronic commerce: theory and implementation. *Decis. Support Syst.*, 24(1):17–27, November 1998.